

Ontolexical resources for feature based opinion mining : a case-study

Anaïs Cadilhac
IRIT
Toulouse University
cadilhac@irit.fr

Farah Benamara
IRIT
Toulouse University
benamara@irit.fr

Nathalie Aussenac-Gilles
IRIT
Toulouse University
ausсенac@irit.fr

Abstract

Opinion mining is a growing research area both at the natural language processing and the information retrieval communities. Companies, politicians, as well as customers need powerful tools to track opinions, sentiments, judgments and beliefs that people may express in blogs, reviews, audios and videos data regarding a product/service/person/organisation/etc. This work describes our contribution to feature based opinion mining where opinions expressed towards each feature of an object or a product are extracted and summarized. The state of the art has shown that the hierarchical organization of features is a key step. In this context, our goal is to study the role of a domain ontology to structure and extract object features as well as to produce a comprehensive summary. This paper presents the developed system and the experiments we carried out on a case study: French restaurant reviews. Our results show that our approach outperforms standard baselines.

1 Introduction

Opinion mining is a growing research area both in natural language processing and information retrieval communities. Companies, politicians, as well as customers need powerful tools to track opinions, sentiments, judgments and beliefs that people may express in blogs, reviews, audios and videos data regarding a product/service/person/organisation/etc. The importance of emotion-oriented computing in the

Web 2.0 has encouraged the creation of new search engines (like Tweetfeel (www.tweetfeel.com)) as well as the creation of a new research group within the W3C, namely the *Emotion Markup Language*, that aims to develop a representation language of the emotional states of a user or the emotional states to be simulated by a user interface. In addition, most information retrieval evaluation campaigns (TREC, NTCI, etc.) have already integrated an opinion track.

Computational approaches to sentiment analysis focus on extracting the affective content of a text from the detection of expressions of “bag of sentiment words” at different levels of granularity. These expressions are assigned a positive or a negative scalar value, representing a positive, a negative or neutral sentiment towards some topic. Roughly, research in this field can be grouped in four main categories (which are not exclusive):

- *Development of linguistic and cognitive models of opinion/sentiment* where already existing psycholinguistic theories of emotions are used to analyse how opinions are lexically expressed in texts (Wiebe et al, 2005; Read et al, 2007; Asher et al, 2009)
- *Elaboration of linguistic resources* where corpus based and dictionary based approaches are used to automatically or semi-automatically extract opinion bearing terms/expressions as well as their sentiment orientation (Strapparava et al., 2004; Turney and Littman, 2002)
- *Opinion extraction/analysis at the document (Pang et al., 2002; Turney, 2002), at the sentence or at the clause level (Kim et al., 2006; Choi et al., 2005) where local*

opinions are aggregated in order to compute the overall orientation of a document/sentence/clause.

- *Feature based opinion mining* (Hu and Liu, 2004; Popescu and Etzioni, 2005; Carenini et al., 2005; Cheng and Xu, 2008) where opinions expressed towards the features of an object or a product are extracted and summarized.

The work described in this paper fits into the last category. The aim is not to compute the general orientation of a document or a sentence, since a positive sentiment towards an object does not imply a positive sentiment towards all the aspects of this object, as in: I like this restaurant even if the service is slow. In feature based opinion mining, a holder (the person who posts the review) expresses a positive/negative or neutral opinions towards a main topic (the object or the product on which the holder expresses his opinions) and its associated features. As defined in (Hu and Liu, 2004), a feature can be a “part-of” of a topic (such as the screen of a camera) or a property of the “part-of” of the topic (such as the size of the screen). The expressed opinion can be explicit, as in “the screen of this camera is great”, or implicit, as in “the camera is heavy”, that expresses a negative opinion towards the weight of the camera. Same features can also be expressed differently, for example, “drink” and “beverage” refer to the same restaurant feature.

Having, for an object/product, the set of its associated features $F=\{f_1,\dots,f_n\}$, research in feature based opinion mining mostly focus on extracting the set F from reviews, and then, for each feature f_i of F , extract the set of its associated opinion expressions $OE=\{OE_1,\dots,OE_j\}$. Once the set of couples (f_i, OE) were extracted, a summary of the review is generally produced. During this process, the key questions are: how the set F of features can be obtained? How they are linguistically expressed? How they are related to each other? Which knowledge representation model can be used to better organize product features and to produce a comprehensive summary?

To answer these questions, we propose in this paper to study the role of an ontology in feature based opinion mining. More precisely, our aim

is to study how a domain ontology can be used to:

- *structure features*: we show that an ontology is more suitable than a simple hierarchy where features are grouped using only the “is-a” relation (Carenini et al., 2005; Blair-Goldensohn et al., 2008)
- *extract explicit and implicit features from texts*: we show how the lexical component as well as the set of properties of the ontology can help to extract, for each feature, the set of the associated opinion expressions.
- *produce a discourse based summary of the review*: we show how the ontology can guide the process of identifying the most relevant discourse relations that may hold between elementary discourse units.

The paper is organised as follows. We give in section 2, a state of the art of the main approaches used in the field as well as the motivations of our work. We present in the next section, our approach. Finally, in section 4, we describe the experiments we carried out on a case study: French restaurant reviews

2 Feature based Opinion mining

2.1 Related Works

Overall, two main families of work stand out: those that extract a simple list of features and those that organize them into a hierarchy using taxonomies or ontologies. The feature extraction process mainly concerns explicit features.

Works without knowledge representation models

The pioneer work in feature based opinion mining is probably the one of Hu and Liu (2004) that applies association rule mining algorithm to discover product features (nouns and noun-phrases). Heuristics (frequency of occurrence, proximity with opinion words, etc...) can eliminate irrelevant candidates. Opinion expressions (only adjective phrases) which are the closest to these features are extracted. A summary is then produced and displays, for each feature, both positive and negative phrases and the total number of these two categories.

To improve the feature extraction phase, Popescu and Etzioni (2005) suggest in their system

OPINE, to extract only nominal groups whose frequency is above a threshold determined experimentally using the calculation of PMI (Point-wise Mutual Information) between each of these nouns and meronymy expressions associated with the product. No summary is produced.

The main limitation of these approaches is that there are a great many extracted features and there is a lack of organization. Thus, similar features are not grouped together (for example, in restaurant domain, “*atmosphere*” and “*ambiance*”), and possible relationships between features of an object are not recognized (for example, “*coffee*” is a specific term for “*drink*”). In addition, polarity analysis (positive, negative or neutral) of the document is done by assigning the dominant polarity of opinion words it contains (usually adjectives), regardless of polarities individually associated to each feature.

Works using feature taxonomies. Following works have a different approach: they do not look for a “basic list” of features but rather a list hierarchically organized through the use of taxonomies. We recall that a taxonomy is a list of terms organized hierarchically through specialization relationship type “is a sort of”. Carenini et al. (2005) use predefined taxonomies and semantic similarity measures to automatically extract classic features of a product and calculate how close to predefined concepts in the taxonomy they are. This is reviewed by the user in order to insert missing concepts in the right place while avoiding duplication. The steps of identifying opinions and their polarity and the production of a summary are not detailed. This method was evaluated on the product review corpus of Hu and Liu (2004) and resulted in a significant reduction in the number of extracted features. However, this method is very dependent on the effectiveness of similarity measures used.

In their system PULSE, Gamon et al. (2005) analyze a large amount of text contained in a database. A taxonomy, including brands and models of cars, is automatically extracted from the database. Coupled with a classification technique, sentences corresponding to each leaf of the taxonomy are extracted. At the end of the process, a summary which can be more or less detailed is produced.

The system described in (Blair-Goldensohn et al., 2008) extracts information about services, aggregates the sentiments expressed on every aspect and produces a summary. The automatic feature extraction combines a dynamic method, where the different aspects of services are the most common nouns, and a static method, where a taxonomy grouping the concepts considered to be the most relevant by the user is used to manually annotate sentences. The results also showed that the use of a hierarchy significantly improves the quality of extracted features.

Works using ontologies. These works aim at organizing features using a more elaborated model of representation: an ontology. Unlike taxonomy, ontology is not restricted to a hierarchical relationship between concepts, but can describe other types of paradigmatic relations such as synonymy, or more complex relationships such as composition relationship or space relationship.

Overall, extracted features correspond exclusively to terms contained in the ontology. The feature extraction phase is guided by a domain ontology, built manually (Zhao and Li, 2009), or semi-automatically (Feiguina, 2006; Cheng and Xu, 2008), which is then enriched by an automatic process of extraction / clustering of terms which corresponds to new feature identification.

To extract terms, Feiguina (2006) uses pattern extraction coupled to a terminology extractor trained over a set of features related to a product and identified manually in a few reviews. Same features are grouped together using semantic similarity measures. The system OMINE (Cheng and Xu, 2008) proposes a mechanism for ontology enrichment using a domain glossary which includes specific terms such as words of jargon, abbreviations and acronyms. Zhao and Li (2009) add to their ontology concepts using a corpus based method: sentences containing a combination of conjunction word and already recognized concept are extracted. This process is repeated iteratively until no new concepts are found.

Ontologies have also been used to support polarity mining. For example, (Chaovalit and Zhou, 2008) manually built an ontology for movie reviews and incorporated it into the polarity clas-

sification task which significantly improve performance over standard baseline.

2.2 Towards an ontology based opinion mining

Most of the researchers actually argue that the use of a hierarchy of features improves the performance of feature based opinion mining systems. However, works that actually use a domain ontology (cf. last section) exploit the ontology as a taxonomy using only the is-a relation between concepts. They do not really use all data stored in an ontology, such as the lexical components and other types of relations. In addition, in our knowledge, no work has investigated the use of an ontology to produce comprehensive summaries.

We think there is still room for improvement in the field of feature based sentiment analysis. To get an accurate appraisal of opinion in texts, it is important for NLP systems to go beyond explicit features and to propose a fine-grained analysis of opinions expressed towards each feature. Our intuition is that the full use of ontology would have several advantages in the domain of opinion mining to:

Structure features: ontologies are tools that provide a lot of semantic information. They help to define concepts, relationships and entities that describe a domain with unlimited number of terms. This set of terms can be a significant and valuable lexical resource for extracting explicit and implicit features. For example, in the following restaurant review: *cold and not tasty* the negative opinion *not tasty* is ambiguous since it is not associated to any lexicalised feature. However, if the term *cold* is stored in the ontology as a lexical realization of the concept *quality of the cuisine*, the opinion *not tasty* can be easily associated to the feature *cuisine* of the restaurant (note that the conjunction *and* plays an important role in the desambiguation process). We discuss this point at the last section of the paper.

Extract features: ontologies provide structure for these features through their concept hierarchy but also their ability to define many relations linking these concepts. This is also a valuable resource for structuring the knowledge obtained during feature extraction task. In addition,

the relations between concepts and lexical information can be used to extract implicit features. For example, if the concept *customer* is linked to the concept *restaurant* by the relation *to eat in*, a positive opinion towards the restaurant can be extracted from the review: *we eat well*. Similarly, if the concept *restaurant* is linked to the concept *landscape* with the relation *to view*, a positive opinion can be extracted towards *the look out of the restaurant* from the following review: *very good restaurant where you can savour excellent Gratin Dauphinois and admire the most beautiful peak of the Pyrénées*

Produce summaries. Finally, we also believe that ontologies can play a fundamental role to produce well organised summary and discursive representation of the review. We further detail this point at the last section of the paper.

3 Our approach

Our feature based opinion mining system needs three basic components: a lexical resource L of opinion expressions, a lexical ontology O where each concept and each property is associated to a set of labels that correspond to their linguistic realizations and a review R.

Following the idea described in (Asher et al, 2009), a review R is composed of a set of elementary discourse units (EDU). Using the discourse theory SDRT (Asher and Lascarides 2003) as our formal framework, an EDU is a clause containing at least one elementary opinion unit (EOU) or a sequence of clauses that together bear a rhetorical relation to a segment expressing an opinion. An EOU is an explicit opinion expression composed of a noun, an adjective or a verb with its possible modifiers (actually negation and adverb) as described in our lexicon L.

We have segmented conjoined NPs or APs into separate clauses—for instance, *the film is beautiful and powerful* is taken to express two segments: *the film is beautiful* and *the film is powerful*. Segments are then connected to each other using a small subset of “veridical” discourse relations, namely:

- **Contrast (a,b),** implies that a and b are both true but there is some defeasible implication

of one that is contradicted by the other. Possible markers can be *although*, *but*.

- *Result(a,b)* indicated by markers like *so*, *as a result*, indicates that the EDU *b* is a consequence or result of the EDU *a*.
- *Continuation(a,b)* corresponds to a series of speeches in which there are no time constraints and where segments form part of a larger thematic. For example, "*The average life expectancy in France is 81 years. In Andorra, it reaches over 83 years. In Swaziland it does not exceed 85 years.*"
- *Elaboration(a,b)* describes global information that was stated previously with more specific information. For example, "*Yesterday, I spent a wonderful day. I lounged in the sun all morning. I ate in a nice little restaurant. Then at night, I met my friend Emily.*"

In a review *R*, an opinion holder *h* comments on a subset *S* of the features of an object/product using some opinion expressions. Each feature corresponds to the set of linguistic realizations of a concept or a property of the domain ontology *O*. For example, in the following product review, EDUs are between square brackets, EOUs are between embraces whereas object features are underlined. There is a contrast relation between the EDU_b and EDU_c which makes up the opinion expressed within the EDU_d.

[I bought the product yesterday] _a, [Even if the product is {excellent}]_b, [the design and the size are {very basic}] _c, [which is {disappointing} in this brand] _c.

The figure below gives an overview of our system. First, each review *R* is parsed using the French syntactic parser Cordial¹, which provides, for each sentence, its POS tagging and the set of dependency relations. The review is then segmented in EDUs using the discourse parser described in (Afantenos and al, 2010).

For each EDU, the system :

1. Extracts EOUs using a rule based approach
2. Extracts features that correspond to the process of term extraction using the domain ontology

¹ http://www.synapse-fr.com/Cordial_Analyseur/




Figure 1 Overview of our system.

3. Associates, for each feature within an EDU, the set of opinion expressions
4. Produces a discourse based summary.

Since the summarization module is not done yet, we detail below the three first steps.

3.1 Extracting Elementary Opinion Units

We recall that an EOU is the smallest opinion unit within an EDU. It is composed of one and only one opinion word (a noun, an adjective or a verb) possibly associated with some modifiers like negation words and adverbs. For example, "*really not good*" is an EOU. An EOU can also be simply an adverb as in *too spicy*. Adverbs are also used to update our opinion lexicon, as in *too chic* where the opinion word *chic* is added. Finally, we also extract expressions of recommendation, such as : *go to this restaurant, you will not regret it*, which are very frequent in reviews.

3.2 Extracting features

This step aims at extracting for the review all the labels of the ontology. Since each concept and its associated lexical realizations correspond to explicit features, we simply project the lexical component of the ontology in the review in order to get, for each EDU, the set of features *F*. Of course, since our lexical ontology does not

cover all the linguistic realizations of concepts and properties in a given domain, many terms in the review can be missed. We show, in the next section, that linking features to opinion expressions can partially solve this problem.

To extract implicit features, ontology properties are used. We recall that these properties define relations between concepts of the ontology. For example, the property “*look at*” links “*customer*” and “*design*” concepts.

3.3 Associating opinions expressions to extracted features

In this step, the extracted opinion expressions in step 1 have to be linked to the features extracted in step 2 i.e. we have to associate to each EDU_i the set of couples (f_i, OE_i). During this step, we distinguish the following cases :

Case 1. Known features and known opinion words. For example, if the lexicon contains the words *really*, *good* and *excellent* and the ontology contains the terms *eating place* and *food* as a linguistic realization of the concepts *restaurant* and *food*, then this step allows the extraction from the EDU “*really good restaurant with excellent food*” the couples (*restaurant*, *really good*) and (*food*, *excellent*). This example is quite simple but in many cases, features and opinion words are not close to each other which make the link difficult to find. Actually, our system deals with conjunctions (including commas) as in: “*I recommend pizzas and ice creams*”, “*very good restaurant but very expensive*”

Case 2. Known features and unknown opinion expressions, as in the EDU “*acceptable prices*” where the opinion word *acceptable* has not been extracted in step 1 (cf. section 3.1). In this case, the opinion lexicon can be automatically updated with the retrieved opinion word.

Case 3. Unknown features and known opinion expressions, as in the EDU “*old fashion restaurant*” where the features *fashion* has not been extracted in step 2 (cf. section 3.2). In this case, the domain ontology can be updated by adding a new label to an existing concept or property or by adding a new concept or a new property in the right place to the ontology. However, since a user may express an opinion on different objects

within a review, this step has to be done carefully. To avoid errors, we propose to manually update the ontology.

Case 4. Opinion expressions alone, as in the EDU “*It’s slow, cold and not good*”. This kind of EDU expresses an implicit feature. In this case, we use the ontology properties in order to retrieve the associated concept in the ontology. For example, in the sentence “*we eat very well*”, the property “*eat*” of the ontology which links “*customer*” and “*food*” will allow the system to determine that “*very well*” refers to “*food*”.

Case 5. Features alone, as in the EDU: “*Nice surrounding on sunny days with terrace*”, even if the feature “*terrace*” is not associated to any opinion word, it is important to extract this information because it gives a positive opinion towards the restaurant. An EDU with features alone can also be an indicator of the presence of an implicit opinion expression towards the feature as in *this restaurant is a nest of tourists*

Actually, our system deals with all these cases except the last one.

4 Case study : mining restaurant reviews

In this section, we present the experiments we carried out on a case study: French restaurant reviews.

4.1 Corpus

For our experiments, we use a corpus of 58 restaurant reviews (40 positive reviews and 18 negatives reviews, for a total of 4000 words) extracted from the web site Qype². Each review contains around 70 words and is composed of free comments on restaurants (but also on other objects like pubs, cinemas, etc.) with a lot of typos and syntactic errors. Each review appears in the web site with additional information such as the date of the review, the user name of the holder and a global rate from 1 (bad review) to 5 (very good review). In this experiment, we only use the textual comments posted. Figure 2 shows an example of a review form our corpus.

² <http://www.qype.fr>

<Start> Un bon petit resto sympa, près du centre, lumineux et à la déco sympa. Le service est de qualité, rapide et on y mange sain et bon. Je recommande ! <Stop>

Figure 2. Example of a restaurant review

4.2 Ontology

Since our aim is to study the role of a domain ontology to feature based opinion mining, we choose to reuse an existing ontology. However, for the restaurant domain, we do not find any public available ontology for French. We thus use a pre-existent ontology³ for English as a basis coupled with additional information that we gather from several web sites⁴. We first translate the existing ontology to French and then adapt it to our application by manually re-organize, add and delete concepts in order to describe important restaurant features. Disparities between our ontology and the one we found in the web mainly come from cultural considerations. For example, we do not found in the English ontology concepts like *terrace*.

Our domain ontology has been implemented under Protégé⁵ and actually contains 239 concepts (from which we have 14 concepts directly related to the superclass owl:think), 36 object properties and 703 labels (646 labels for concepts and 57 labels for properties). The left part of figure 3 shows an extract of our restaurant domain ontology.

4.3 Opinion Lexicon

Our lexicon contains a list of opinion terms where each lexical entry is of the form:

[POS, opinion category, polarity, strength] where *POS* is the part of speech tagging of the term, *opinion category* can be a judgment, a sentiment or an advice (see (Asher et al, 2009) for a detailed description of these categories), *polarity* and *strength* corresponds respectively to the opinion orientation (positive, negative and neutral) and the opinion strength (a score between 0 and 2). For example, we have the following entry for the term *good*: [Adj, judgment, +, 1].

³ <http://gaia.fdi.ucm.es/ontologies/restaurant.owl>

⁴ <http://www.kelrestaurant.com/dept/31/> and <http://www.resto.fr/default.cfm>

⁵ <http://protege.stanford.edu/>

The lexicon actually contains 222 adjectives, 152 nouns, 157 verbs. It is automatically built following the algorithm described in (Chardon, 2010). We then add manually to this lexicon 98 adverbs and 15 expressions of negation.

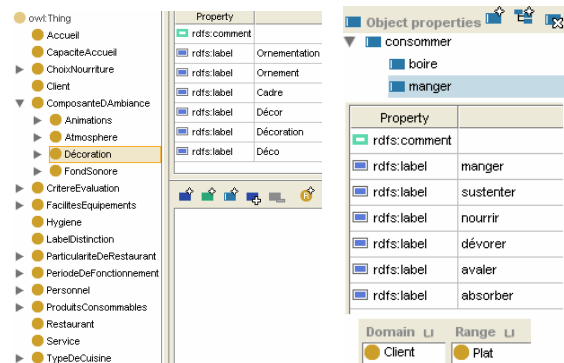


Figure 3. Extract of the restaurant domain ontology : Left - hierarchy of concepts and labels of “decoration” concept. Right – information about a particular object property.

4.4 Experiments

We conduct three types of experiment: the evaluation of the extraction of elementary opinion units (cf. section 3.1), the evaluation of the features extraction step (cf. section 3.2) and finally, the evaluation of the link between the retrieved opinion expressions and the retrieved object features (cf. section 3.3).

These experiments are carried out using GATE⁶ toolkit. To evaluate our system, we create a gold standard by manually annotate in the corpus implicit and explicit elementary opinion units, implicit and explicit object features as well as for each opinion expression its associated feature.

Evaluation of the EOU extraction step.

The table below shows our results. Our system misses some EOU for two main reasons. The first one is due to missed opinion words in the lexicon and to implicit opinion expressions, such as *breathhtaking*, since our extraction rules do not manage these cases (note that implicit opinion detection is still an open research problem in opinion mining).

⁶ <http://gate.ac.uk/>

The second reason is the errors that come from the syntactic parser mainly because of typos and dependency link errors. Concerning precision, false positives are mainly due to some opinion words that are in our lexicon but they do not express opinions in the restaurant domain. In addition, some of our extraction rules, especially those that extract expression of recommendations, do not perform very well which imply a loss of precision.

Precision	0,7486
Recall	0,8535
F-measure	0,7976

Table 1. Evaluation of EOU extraction

Evaluation of the features extraction step.

Since the corpus is in the restaurant domain, the precision of this task is very good because most of the extracted features are relevant. However, recall is not as good as precision because the set of ontology labels do not totally cover the terms of the corpus. Another limitation of our system is that we do not take into account the cases where a term can be a linguistic realization of many concepts (ex. café can be a drink or a place to drink).

Figure 4 shows an example of the result we obtain for this step.

<Start> Un bon petit resto sympa, près du centre, lumineux et à la déco sympa. Le service est de qualité, rapide et on y mange sain et bon. Je recommande ! <Stop>

Figure 4. Result of EOU (blue) and ontological term (pink) extraction

Evaluation of the link between EOU and features.

The figure below shows our result on a sample. In this example, the system is able to extract opinion expressions which do not contain words present in the lexicon. It is the case with “*sympa (nice)*” which has been correctly associated to “*resto (restaurant)*” and “*deco (interior design)*” even if the word *nice* was not in the lexicon.

In order to evaluate the added value of using an ontology to feature based opinion mining, we compare our system to the well known ap-

proaches of Hu and Liu and Popescu and Etzioni (cf. section 2.1) that do not use any knowledge representation. We have also compared our approach to those that use taxonomies of concepts by deleting the properties of our domain ontology. The results are shown in table 2.

<Start> Un bon petit resto sympa, près du centre, lumineux et à la déco sympa. Le service est de qualité, rapide et on y mange sain et bon. Je recommande ! <Stop>

Figure 5. Result of linking EOU to extracted features

	Precision	Recall	F-measure
Our system	0,7692	0,7733	0,7712
Hu and Liu	0,6737	0,7653	0,7166
Popescu and al	0,7328	0,7387	0,7357
Taxonomy	0,7717	0,7573	0,7644

Table 2. Evaluation of our system and its comparison to existing approaches

In the Hu and Liu approach, features are nominal groups. We first extract all frequent features from our corpus that appear in more than 1% of the sentences. Then we extract EOU from those sentences (note that contrary to Hu and Liu, we do not extract only adjectives, but also nouns, verbs and adverbs). Non frequent features are finally removed as described in (Hu and Liu, 2004). In order to improve the extraction of relevant features, we extract features that have a good point mutual information value with the word *restaurant*, as described in (Popescu and Etzioni, 2005). The precision of our system is better compared to the approach of Hu and Liu that extracts too many irrelevant features (such as *any doubt*, *whole word*). Our system is also better compared to the PMI approach even if it performs better than Hu and Liu’s approach. Recall is also better because our system can extract implicit features such as *well eating*, *lot of noise*, thanks to the use of ontology properties. Finally, when using only taxonomy of concepts instead of the ontology, we observe that the F-measure is slightly better because actually fea-

tures related to object properties represent only 1,6% of feature cases in our corpus. Using, the ontology, our approach is able to extract from sentences like "we eat good and healthy" the couples (eat, good) and (eat, healthy) and then to link the opinion expressions to the concept *dish* whereas when using only the taxonomy, these opinion expressions are related to any feature.

5 Conclusion and prospects

5.1 Contribution of our system

Our method is promising because the use of the ontology allows to improve the feature extraction and the association between an opinion expressions and object features. On the one hand, the ontology is useful thanks to its concept list which brings a lot of semantic data in the system. Using concept labels the ontology allows to recognize terms which refer to the same concepts and brings some hierarchy between these concepts. On the other hand, the ontology is useful thanks to its list of properties between concepts which allows recognizing some opinions expressed about implicit features.

5.2 Prospects

Opinion lexicon improvement.

The opinion extraction we achieved is naive because we use a simple opinion word lexicon which is not perfectly adapted to the domain. To improve this part of the treatment, it would be interesting to use opinion ontology. As illustrated in section 2.2, constructing a domain ontology for the purpose of opinion mining poses several interesting questions in term of knowledge representation, such as: what are the frontiers between knowledge, where concepts are domain dependent, and opinion, where expressions can be at the same time dependent (the term *long* can be positive for a *battery life* but negative if it refers to a the *service* of a restaurant) and independent (the term *good* is positive) from a domain. Our intuition is that the two levels have to be separated as possible.

Natural Language processing (NLP) rules improvement.

Our system is limited by some current NLP problems. For example, the system does not

treat the anaphora. For example, in the sentence "*Pizzas are great. They are tasty, original and generous*", it does not recognize that the three last adjectives refer to "*pizzas*". There is also the problem of conditional proposition. For example, in the sentence "*affordable prices if you have a fat wallet*", the system is not able to determine that "*affordable prices*" is subject to a condition.

Ontology and lexicon enrichment.

Thanks to the ability to link opinion expression and ontological term extractions, our system is able to extract some missing opinion words and labels of the ontology. We think it could be interesting to implement a module which allows the user to easily enrich opinion word lexicon and ontology. Furthermore, it will be interesting to evaluate the benefit of this method in both opinion mining and ontological domains.

Towards a discourse based summary.

The last step of the system is to produce a summary of the review that presents to the user all the opinion expressions associated to the main topic and all its features. This summary does not pretend to aggregate opinions for each feature or for the global topic. Instead, the aim is to organize the opinions of several reviews about one restaurant in order to allow the user to choose what feature is important or not for him. In addition to this kind of summarization, we want to investigate how the domain ontology can be used to guide the process of identifying the most relevant discourse relations between elementary discourse units (EDU). Actually, the automatic identification of discourse relations that hold between EDUs is still an open research problem. Our idea is that there is continuation relation between EDU that contain terms that refer to concepts which are at the same level of the ontology hierarchy, and there is an elaboration relation when EDU contains more specific concepts than those of the previous clause.

References

Afantenos Stergos, Denis Pascal, Muller Philippe, Danlos Laurence. *Learning Recursive Segments for Discourse Parsing*. LREC 2010

- Asher, Nicholas, Farah Benamara, and Yvette Y. Mathieu. 2009. *Appraisal of Opinion Expressions in Discourse*. *Linguisticae Investigationes*, John Benjamins Publishing Company, Amsterdam, Vol. 32:2.
- Asher Nicholas and Lascarides Alex. *Logics of Conversation*. Cambridge University Press, 2003
- BlairGoldensohn, Sasha, Kerry Hannan, Ryan McDonald, Tyler Neylon, George A. Reis, and Jeff Reynar. 2008. *Building a Sentiment Summarizer for Local Service Reviews*. WWW2008 Workshop: Natural Language Processing Challenges in the Information Explosion Era (NLPIX 2008).
- Carenini, Giuseppe, Raymond T. Ng, and Ed Zwart. 2005. *Extracting Knowledge from Evaluative Text*. In Proceedings of the 3rd international conference on Knowledge capture.
- Chardon Baptiste. *Catégorisation automatique d'adjectifs d'opinion à partir d'une ressource linguistique générique*. In proceedings of RECITAL 2010, Montreal, Canada
- Pimwadee Chaovalit, Lina Zhou: *Movie Review Mining: a Comparison between Supervised and Unsupervised Classification Approaches*. HICSS 2005
- Cheng, Xiwen, and Feiyu Xu. 2008. *Fine-grained Opinion Topic and Polarity Identification*. In Proceedings of the Sixth International Language Resources and Evaluation (LREC' 08), Marrakech, Morocco.
- Feiguina, Olga. 2006. *Résumé automatique des commentaires de Consommateurs*. Mémoire présenté à la Faculté des études supérieures en vue de l'obtention du grade de M.Sc. en informatique, Département d'informatique et de recherche opérationnelle, Université de Montréal.
- Gamon, Michael, Anthony Aue, Simon Corston-Oliver, and Eric Ringger. 2005. *Pulse: Mining Customer Opinions from Free Text*. In Proceedings of International symposium on intelligent data analysis N°6, Madrid.
- Hu, Minqing, and Bing Liu. 2004. *Mining and Summarizing Customer Reviews*. In Proceedings of the 10th ACM SIGKDD international conference on Knowledge discovery and data mining.
- Kim, Soo-Min, and Eduard Hovy. 2006. *Extracting Opinions, Opinion Holders, and Topics Expressed in Online News Media Text*. In Proceedings of ACL/COLING Workshop on Sentiment and Subjectivity in Text, Sydney, Australia.
- Pang, Bo, Lillian Lee, and Shivakumar Vaithyanathan. 2002. *Thumbs up? Sentiment Classification using Machine Learning Techniques*. Proceedings of EMNLP 2002.
- Popescu, Ana-Maria, and Oren Etzioni. 2005. *Extracting Product Features and Opinions from Reviews*. In Proceedings of the conference on Human Language Technology and Empirical Methods in Natural Language Processing.
- Read, Jonathon, David Hope, and John Carroll. 2007. *Annotating Expressions of Appraisal in English*. The Linguistic Annotation Workshop, ACL 2007.
- Strapparava, Carlo, and Alessandro Valitutti. 2004. *WordNet-Affect: an Affective Extension of WordNet*. Proceedings of LREC 04.
- Turney, Peter D. 2002. *Thumbs Up or Thumbs Down? Semantic Orientation Applied to Unsupervised Classification of Reviews*. Proceedings of 2006 International Conference on Intelligent User Interfaces (IUI06).
- Turney, Peter D., and Michael L. Littman. 2002. *Unsupervised Learning of Semantic Orientation from a Hundred-Billion-Word Corpus*. National Research Council, Institute for Information Technology, Technical Report ERB-1094. (NRC #44929)
- Wiebe, Janyce, Theresa Wilson, and Claire Cardie. 2005. *Annotating Expressions of Opinions and Emotions in Language*. *Language Resources and Evaluation* 1(2).
- Zhao, Lili, and Chunping Li. 2009. *Ontology Based Opinion Mining for Movie Reviews*. In Proceedings of the 3rd International Conference on Knowledge Science, Engineering and Management.